

Introduction to Issues in Developing Adaptive Intelligence

Gary Berg-Cross

What are the issues to developing and measuring the intelligence of adaptive systems, how are they defined and what is the background to the issues? Given these questions it seemed to make a sense to present a top-down introduction to this in order to facilitate discussion such as we might have in a panel.

A major context for a track on developing adaptive intelligence includes the foundational assumptions laid out by Albus (1999) in “The Engineering of Mind” which lays out a reference model architecture that serves as a principle vein of work on the design of intelligent systems. Jim Albus (1999) is based on functional, control and goal-oriented pragmatic principles using 2 assumptions. His first assumption is that the function of the mind/brain is to generate and control intelligent behavior. A second assumption is that intelligent behavior is understood as “appropriate action” in an uncertain environment where appropriate action is understood as action which increases the probability of success in achieving high priority goals. While Albus’ approach includes pragmatic concepts it grows out of control theory and what Haugeland calls GOFAI, Good Old Fashion AI”. To this meta model Albus adds a specific architecture for intelligent behavior as the result of goals and plans interacting at many hierarchical levels with knowledge represented in a multi-resolutional world model (discussed at greater length in Albus and Meystel, 2002). A general theoretical reference model for mind is axiomatically defined by Albus using 5 key concepts and their assumptions enables the approach. Briefly, these are function components, structural storage, interconnected computational modules in a control system architecture, hierarchical layering of components and attentional selection. These are depicted in Figure 1 and detailed as follows:

1. The functional elements of an intelligent system are: behavior generation, sensory perception (filter, detect, recognize, interpret), world modeling (to store knowledge, predict and simulate future), and value judgment (to compute cost, benefit, and uncertainty attributes). These elements are detailed by sub-activities. As an example behavior generation is based on planning and control of actions designed to achieve behavioral goals. The planning a process itself is complex and:

1. assigns responsibility to agents/ computational elements for jobs, and allocates resources to agents to perform assigned jobs,
2. hypothesizes strings of actions (plans) for agents from a “vocabulary” of possible actions to accomplish jobs,
3. simulates and predicts the results of executing these hypothesized plans,
4. evaluates the predicted results of the hypothesized plans,
5. selects the hypothesized plan with the most favorable results for execution.

In general Albus’ functional elements, while computational in principle, evoke many appealing cognitive concepts such as meaningful transformations, risk, benefit, importance, advantageous goals etc. Biological correlates to functions are noted (but since this is a moving target, may need to be revised over time).

2. These functional elements of an intelligent system are supported by a knowledge database that stores (both long-term and short-term) a priori and dynamic information about the world in the form of state variables to record state-of-the-world, symbolic entities (as in GOFAI), symbolic events, rules and equations, structural and dynamic models, task knowledge, signals, images, and maps. It is worth noting in passing that Joslyn (2000) understands Albus’s overall approach to be one of semiotic control system using models, such as stored in a

knowledge base providing the special property of intelligence and combined with control systems to purposely organize functions such as listed above.

3. In the reference architecture the functional elements and knowledge database described above can be implemented by a set of computational modules that are interconnected to make up nodes in a control system architecture. Exemplars of such nodes are found in the Real-time Control System (RCS) built at NIST. Each node is part of a control system that performs the 4 basic processes -sensing, maintains a world model, computing values, and generating behavior and their supporting sub-processes (planning, task decomposition etc.). To Albus a node corresponds to a functional set of brain neurons closing the loop between afferent and efferent neural pathways. Interestingly conceptually this architecture does not directly address components that do not loop between afferent and efferent parts and thus makes them a secondary sub-component.
4. The complexity inherent in intelligent systems can be managed through hierarchical layering, which is a common method for organizing complex systems that has been used in many different types of organizations throughout history for effectiveness and efficiency of command and control. A key to hierarchical control is that higher level nodes have broader scope and longer time horizons, with less concern for detail, while lower level nodes have narrower scope and shorter time horizons, with more focus on detail. Elements may be deliberative or reflexive throughout the hierarchy
5. The complexity of the real world environment can be managed through a strategic process of focusing attention. Since the world is not uniformly important this serves the needs of an agent with limited computational resources.
Attention focuses sensors on task-relevant objects.

Taken as a whole Albus believes that work based on this reference model helps advance the scientific inquiry into of the nature of mind, and also will very likely also lead to practical improvements in intelligent machine systems technology for many fields. A number of practical achievements can be pointed to, but it may be useful to relook at the overall architecture, its assumptions and details.

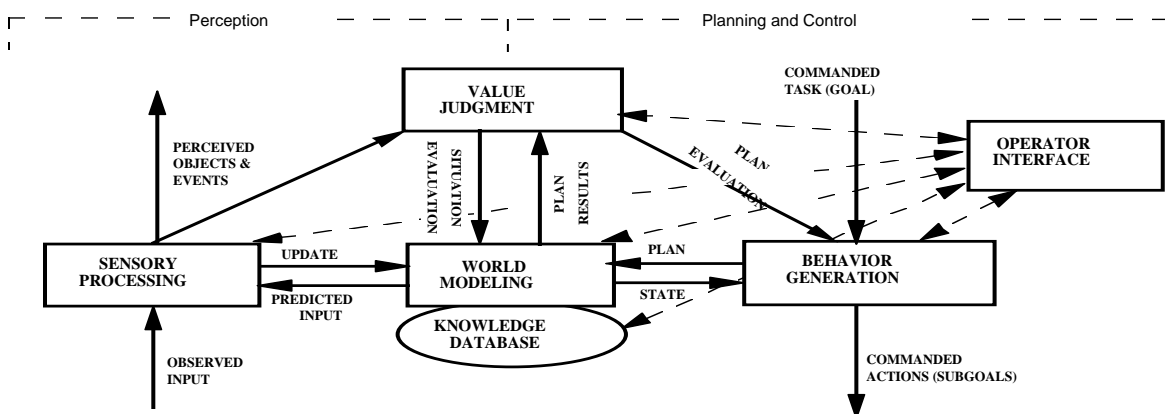


Figure 1 RELATIONSHIPS BETWEEN ALBUS' FUNCTIONAL ELEMENTS

A context for this discussion is a section from last year's PerMIS conference (Machine Intelligence: Measures & Issues II) which raised some issues relevant to the architecture and its assumptions. These included:

- a. Going beyond sequential, functional, hierarchical architectural ideas to include how intentions and adaptive abilities emerge to provide the roots of "intelligent" agent actions (Berg-Cross, 2003).
- b. How to understand intelligence performance when we consider how natural agents (individual humans and animals) readily adapt to rapidly changing and/or complex environments including rare, ambiguous, and unexpected events. An example of how might artificial intelligence might adapt was presented by Gunderson (2003) considering a robotic intelligence that has a robust ability to "develop" hypotheses to solve problems and thus starts chasing the family dog to more effectively (intelligently) perform housekeeping.
- c. Moving towards "natural intelligence" might intelligent information-processing bridge between differently-scaled models by means of cross-scalar coherence and autonomy-negotiation to create a hyperscalar system (Cottam, 2003).
- d. How valuable are functions of creativity and playfulness in intelligent systems when considered as part of an adaptive learning and developmental process? Is there a common architecture that includes interaction, adaptation, innovation and "immunity" (Arata, 2003).
- e. Strong AI vs. Weak AI asked the question on how important specific models found in knowledge bases versus general and abstract know. But much of the human knowledge seems inconsistent, loosely organized, and in perpetual flux rather than faithful copies of reality found in implementations of both strong or weak AI. The mind may neither be a highly organized knowledge base nor a large set of fuzzy images. A concept developed by John Sowa is *knowledge soup*. Knowledge is fluid, lumpy, with adherable chunks of Theories, Models and Hypotheses that float in and out of awareness. Is it useful to of considering agent knowledge as a fluid rather than precise system of "true" facts (Berg-Cross, 2003).

Taken together these question some of the foundational assumptions laid out by Albus in "The Engineering of Mind" and suggests modifications if not major overhaul to reference model architecture. We agree that the scope to resolve these issues will be broad requiring understanding and cooperative work from:

1. Neurosciences
2. Cognitive Sciences and cognitive information processing
3. Artificial Intelligence
4. Learning and Complex Adaptive Systems
5. Intelligent Control
6. Robotics and Intelligent Machines
7. Intelligent Manufacturing Systems
8. Game Theory and Operations Research
9. Image Understanding
10. Planning and Reasoning
11. Philosophy
12. Linguistics and Speech Understanding

As an example, much of neuroscience models (for example perceptual cognition) are network rather than hierarchical models. While it is good for engineering to reuse the same functional the basic "unit" of neural computation does not seem to share the same processing components as we move from one level to another.

Berg-Cross (2003) notes separate, but fundamentally different functioning to support the high-level belief-intention parts of intelligent behavior and, at least in human architectures, these seem significantly different from the lower level functions. This is related to the argument of how easy or practical it currently is to engineer intelligent systems (hierarchy simplifies the engineering) as opposed to letting them develop.

Another issue concerns the “meaning” and measurement of autonomy of a system that is learning, modifying its knowledge and adapting its behavior. Understanding the interaction of these components over time is a central part of understanding how an agent arrives at “appropriate action” to increase the probability of success in achieving high priority goals. How important is self-organization and development?

Finally it is worth noting that some “general principles” for complex systems may be useful to discuss for intelligence. The phenomena of pattern formation and self-organization found in nonequilibrium physical, chemical and biological systems may be governed by a number of general principles. This idea, arising in the natural sciences study of complex systems structural development, has now been applied to discussion of the development of intelligence. It is a new way to bridge the gap between what individual element function, such as RCS nodes, and what many of them do when they adapt/learn/evolve to function “cooperatively”.

An issue then is how GOF AI and engineering approaches might be supplemented by research strategies and computational tools growing out of nonlinear dynamical systems. Numerical and analytical investigations of these systems lead to new mathematical results and problems, as well as to formal bridges to other biological and physical systems, notably dissipative systems that describe aspects of self-organization and nonequilibrium behavior. These formal investigations have already suggested new designs for computer vision, adaptive pattern recognition machines, and autonomous robots, and as an integrative approach may provide basic science with designs of adaptive intelligent systems. As an example, in the study of complex systems, especially natural ones, the idea of emergence is used to explain development of patterns, structures, and/or properties that do not seem reducible to a system's existing components and their linear interaction. Emergence becomes of increasing importance as explanatory constructs for complex systems characterized when:

- Global, system organization appears to be of a different kind than functional components alone;
- Components can be changed or even removed without an accompanying loss of function of the higher functions of a system.

In concluding, adaptation and learning are seen as one of the central points of such IS research, but diverse approaches may be meaningfully integrated into thinking about the future of Intelligence Systems. We advocate an integrative approach - let us not be stymied about differences and together keep doing interesting things. In this panel, and in the track as a whole we hope to discuss core ideas and explore approaches that will possibly shape future work in the area of adaptive intelligence.

References

Albus, James.S. (1991). “Outline for a Theory of Intelligence”. *IEEE Transactions on Systems, Man and Cybernetics*, **21** (3): p. 473-509

Albus, James S: (1999) "*The Engineering of Mind*", Information Sciences,v.117, pp. 1-18

Arata, L., **Interactive Measures and Innovation**, presented at PerMIS 2003.

Berg-Cross, G., **A Pragmatic Approach to Discussing Intelligence in Systems**, presented at PerMIS 2003.

Cottam, R., Ranson, W., Vounckx, R., **Abstract or Die: Life, Artificial Life and (v)Organisms**, presented at PerMIS 2003.

Gunderson, J., Gunderson, L., **Mom! The Vacuum Cleaner is Chasing the Dog Again!**, presented at PerMIS 2003.

Joslyn, Cliff: (2000) "**Towards Measures of Intelligence Based on Semiotic Control**", *2000 Workshop on Performance Metrics in Intelligent Systems*, ed. A. Meystel, NIST

Meystel, Alexander M and Albus , James, **Intelligent Systems: Architecture, Design, and Control**
2001